

1
2
3
4
5
6
7
8
9
10

MATHEMATICS AT THE RAND CORPORATION

T. A. Brown

March 1967

D D C
REF ID: A65000000000000000000
APR 18 1967
R
C

P-3543

ARCHIVE COPY

MATHEMATICS AT THE RAND CORPORATION

T. A. Brown*

Many mathematicians work at The RAND Corporation. Of 509 professional workers, 14 percent are mathematicians; 9 percent, programmers; 4 percent operations analysts; 12 percent, physicists; 16 percent economists; 28 percent, engineers; 6 percent, political scientists; and 11 percent, other (the other category includes anthropologists, psychologists, linguists, and the like). It is clear that mathematicians or people with substantial mathematical training constitute a large portion of the staff of The RAND Corporation. Even of the mathematicians, however, perhaps two-thirds are not doing mathematics in the sense of preparing papers which appear in mathematical journals. They are working on problems in military affairs or economics or social science in which there is a substantial mathematical component.

Some feel that if people trained in mathematics are not producing mathematics they are wasting their talents.

* Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This is the text of a talk given to a group of high school mathematics teachers at the Occidental Mathematics Field Day on February 18, 1967.

When I hear such remarks, I am reminded of J. P. Morgan, who was the dominant figure in American high finance during the last half of the 19th century. It is not widely known that J. P. Morgan was trained as a mathematician at the University of Göttingen in Germany. In 1857, at the time of his graduation, the university's noted professor of mathematics, Professor Ulrich, took him aside and had a heart-to-heart talk with him. Professor Ulrich asked about his plans in life, and Pierpont answered that his father intended that he should go into business in America. The professor said, with real feeling, that he hated to see Pierpont leave and that it had been a great satisfaction to teach him. Under the circumstances, he felt constrained to tell him that no matter what plans his father had, Pierpont was making a great mistake to go into business; he really should make mathematics his life work, since he had shown such unusual aptitude in the field. Professor Ulrich's confidence in the young man's mathematical abilities was great; and he assured Pierpont that he could promise after but one year more a position at Göttingen as Instructor in Mathematics and as his assistant. He said that he would not go so far as to lead Pierpont to hope that he would one day succeed Ulrich as Professor of Mathematics—that would be too much to hold out because Professor Ulrich did not have the power to choose his successor. However, he could promise, if Pierpont worked hard and showed proficiency

in teaching, he would use his best efforts to have him chosen as Professor when he himself should get so old that he had to retire. Both Professor Ulrich and Morgan recounted this incident years afterwards, and to Pierpont Morgan it was a cause of great satisfaction.

Was Pierpont's training in mathematics wasted because he became the dominating figure in American economic life in the last half of the 19th century instead of pursuing his mathematical career? I don't think so. I believe that the training he received in mathematics helped him to separate fundamental issues from extraneous issues in economic and financial situations and enabled him to be a better banker and a better economist than he would otherwise have been.

The subject of this talk is not mathematicians at The RAND Corporation, but mathematics at The RAND Corporation. I am going to speak of the work of the one-third of the mathematicians at The RAND Corporation who are doing creative mathematics in the sense of mathematics that is published in mathematical journals and contributes, in general, to the mathematical life of this country and the world. I am going to speak in particular of four categories of mathematical work: linear programming, dynamic programming, flows in networks, and game theory.* With more time

*The RAND studies cited in the Bibliography present comprehensive discussions of these mathematical fields.

I might discuss branching processes, automatic categorization, radiative transfer, invariant imbedding, integer programming, potential theory and markov chains, or any of a number of other areas. But I hope that the four examples I will present suggest at least the flavor of the work which is done by the creative mathematicians at RAND. To each of these areas RAND mathematicians have made substantial contributions.

LINEAR PROGRAMMING

Let us consider a simple example showing what linear programming encompasses. Suppose that R is a real-valued function of two variables, which is convex; that is, R has the property that if t is a number between 0 and 1, then

$$R(tx_1 + (1-t)y_1, tx_2 + (1-t)y_2) \leq t \cdot R(x_1, x_2) + (1-t)R(y_1, y_2).$$

Any linear function of two variables is convex. The square root of $(x_1^2 + x_2^2)$ is convex and it is very easy to think of other examples of convex functions. Suppose that we are seeking to maximize $R(x_1, x_2)$ subject to linear constraints. For the sake of definiteness, let us say that one constraint is that x_1 must be positive; another, that x_2 must be positive; another, that $2x_1 + x_2 \leq 1$; and still another, that $x_1 + 2x_2 \leq 1$. What do these constraints mean?

Consider the (x_1, x_2) -plane (Fig. 1), which I am sure is familiar to all of you. The condition that x_1 be greater than zero means that the part of the plane to the left of the x_2 -axis must be ruled out of consideration in a search for a maximum. Similarly, the condition that x_2 be greater than zero means that the part of the plane

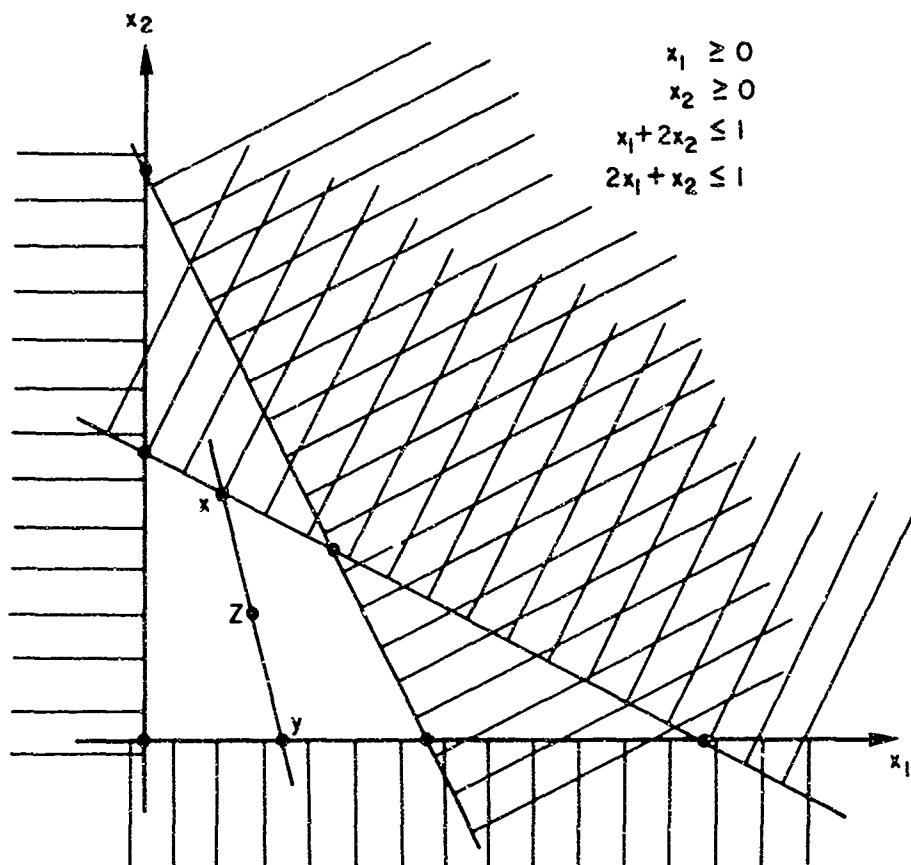


Fig. 1 - An example of a linear programming problem

below the x_1 -axis must be ruled out of consideration. The condition that $2x_1 + x_2 \leq 1$ means that the portion of the plane above the line passing through the point $(0, 1)$ and the point $(1/2, 0)$ must be ruled out of consideration. The final inequality means that the portion of the plane falling above the line passing through $(0, 1/2)$ and $(1, 0)$ must be ruled out of consideration. In other words, we are seeking to find the maximum that this convex function assumes in the convex domain shown in Fig. 1.

The condition that this function be convex means that the maximum must be achieved on one of the vertices of this convex region. To illustrate, suppose the maximum were achieved at some point z in the interior of this region. Take a line passing through z and intersecting the boundary of the region at points x and y . Now z would be a weighted average of x and y —that is, we would have $z = tx + (1 - t)y$ for some $0 < t < 1$. Because of the convexity condition, the value of the function at z must be less than the weighted average of the value of the function at x and the value of the function at y .

Therefore, the value of the function at x or the value of the function at y must be at least as great as it is at z . We see that the maximum must be achieved somewhere on the boundary. If we take a boundary point, a similar argument shows that the maximum must be achieved at a vertex. With these simple inequalities in two variables,

it is easy to solve for these vertices. Solving the first two equations gives us $(0, 0)$ as a vertex. Solving the first and third equations gives us $x_2 = 1/2$; i.e., the point $(0, 1/2)$. Solving the first and fourth equations gives us the point $(0, 1)$. Solving the other three pairs of equations gives us the vertices $(1/2, 0)$, $(1, 0)$, and $(1/3, 1/3)$.

We need not consider the vertices $(0, 1)$ and $(1, 0)$ because they are not feasible—they violate one of our inequalities. We know that the maximum value of R must be achieved at one of the four remaining vertices. We evaluate R at these four points, and whichever value is the greatest would be the maximum value which the function achieves on this convex domain.

This is a very simple process when one is dealing with two variables and with a small number of inequalities. But if we had ten variables and twenty inequalities, there would be as many solutions to consider as there are different ways of selecting ten items out of twenty; that is, the more than 22 million ways of picking ten inequalities out of twenty, which is clearly impractical. This is not an unrealistic problem; in logistics and in chemical kinetics, problems often involve more than ten inequalities and twenty variables. How can you tackle this kind of problem? How can you solve it expeditiously? The answer is that you do not work with inequalities first, then look at the

objective function as an afterthought. Rather, you let the way you look at the inequalities be guided by the objective function. You first find a feasible vertex (i.e., one that satisfies these inequalities). Then you consider (in this case) the ten rays emanating from that point; i.e., the 10 one-dimensional simplices obtained by eliminating one of the ten equalities that define the vertex. You find one of those rays which results in a reduction of the objective function. You follow that ray until you reach another vertex which is on the boundary of the convex domain of feasible points and repeat the process. It is not necessary to solve all possible subsets of ten inequalities from the set of twenty. Only a subset of them need to be solved—in a case like this, perhaps forty or fifty. This would be a hard chore by hand, but it is easy to do on a digital computer.

DYNAMIC PROGRAMMING

Another technique is known as dynamic programming. Let us assume that we have a function R of ten variables, which has the following form:

$$R(x_1, x_2, \dots, x_{10}) = g_1(x_1) + g_2(x_2) + \dots + g_{10}(x_{10});$$

and that we wish to find the maximum of this function subject to the constraint that the sum of the x_i is equal to some fixed value which we will call x . Let us say also that

the x_i must all be nonnegative. At first glance, this may look like just another linear programming problem because we are asked to maximize a given function subject to linear constraints. However, it is not a linear programming problem because we are not assuming that this objective function R is necessarily convex. If R were convex, this would be a very trivial problem. One would solve it by simply looking at the ten points— $(x, 0, 0, \dots, 0)$, $(0, x, 0, \dots, 0)$, $(0, 0, x, \dots, 0)$, and so forth—which are the vertices of the convex set defined by the equality $\sum x_i = x$ and the inequalities $x_1 \geq 0, x_2 \geq 0, \dots, x_{10} \geq 0$. The maximum of this function must be achieved at one of these points. So, in a special case where R is a convex function, the problem is trivial. However, if the function is not convex, the problem is not at all trivial, because the maximum may be achieved at some interior point, at some point inside the area defined by this equality and these inequalities.

There are many traditional approaches to the problem. If these functions are differentiable, one could try calculus. This would lead to a system of equations with the following form:

$$g_1'(x_1) - g_{10}'\left(x - \sum_{i=1}^9 x_i\right) = 0,$$

$$g_2'(x_1) - g_{10}'\left(x - \sum_{i=1}^9 x_i\right) = 0,$$

$$g_3'(x_1) - g_{10}'\left(x - \sum_{i=1}^9 x_i\right) = 0,$$

and so on—a system of nine simultaneous nonlinear equations in nine unknowns. This is not, in general, a workable approach.

A second approach is straight enumeration. We cover this set with a mesh, evaluate the function at each point in that mesh (hopefully, fine enough that we really come close to the maximum), and simply pick the point at which the function achieves its maximum. If the interval from 0 to x is divided into 100 increments, we must consider approximately 3×10^{11} points— 10^{18} divided by 10 factorial. If a computer is able to evaluate this function a thousand times a second, a solution will be achieved by this method in about ten years.

A third approach, the dynamic programming approach, is to exploit the particular structure of the objective function. A set of auxiliary functions f is defined as follows: Let f_1 of y be simply g_1 of y . Let f_2 of y be defined as the maximum, as x_2 ranges between 0 and y , of $g_2(x_2) + f_1(y - x_2)$. In other words, f_2 is the maximum sum of $g_1 + g_2$ which is achieved by allocating y units

between the g_1 and the g_2 part of the process. Similarly, f_3 of y is defined as the maximum for x_3 between 0 and y of $g_3(x_3) + f_2(y - x_3)$, and so on, to f_{10} of y which is defined as the maximum for x_{10} between 0 and y of $g_{10}(x_{10}) + f_9(y - x_{10})$.

Consider what this process involves. Evaluating f_1 at 100 points between 0 and x costs 10^2 evaluations.

To evaluate f_2 at 100 points between 0 and x , for each value y it will be necessary to run from 0 to y in appropriate increments, and on the average there will be about 50 values of x_2 . For the 100 values of y , there will be 50 times 10^2 evaluations to evaluate f_2 and so forth for the other f 's. Altogether there will be 451 times 10^2 evaluations, and a computer that can make a thousand evaluations per second will produce the answer in about 45 seconds.

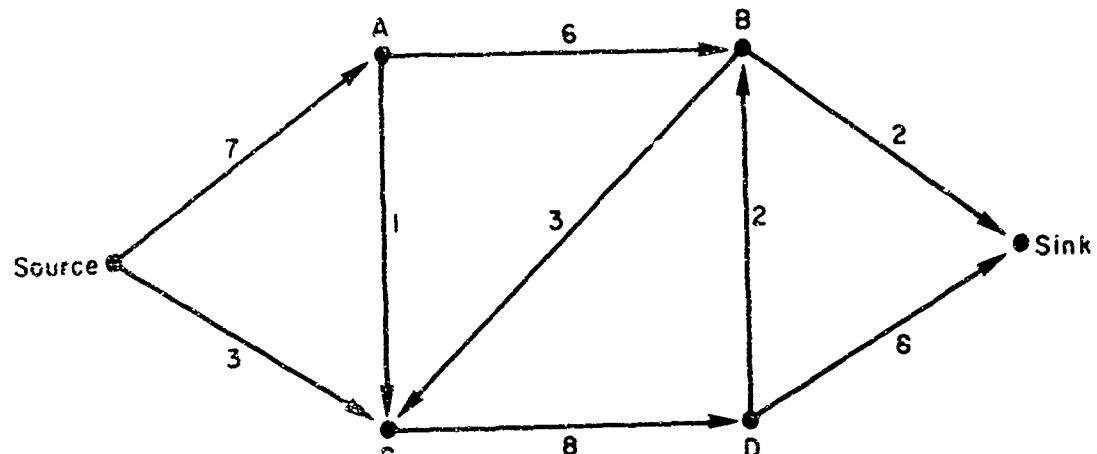
There is more to dynamic programming than we have seen here. The objective function does not always appear in this form. Manipulation of the problem into this form often requires much work and insight. However, I think this discussion gives you the basic idea, and I think it is remarkable that such a simple change in point of view can result in such a tremendous economy in operation.

FLows IN NETWORKS

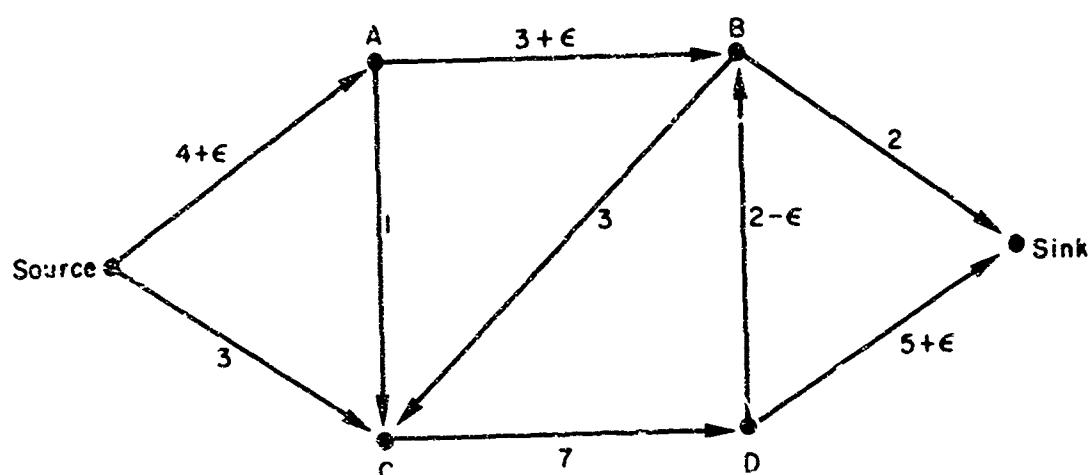
The field of flows in networks was originally regarded as a special case of linear programming; but as time went

on, it was more and more recognized as a separate category in its own right. Figure 2 illustrates a problem of maximizing flow in a simple network. In the network we have a source and a sink, intermittent nodes (A, B, C, D) between the source and the sink, and directed arcs connecting the source to the intermittent nodes and eventually to the sink. Flow over directed arcs is in one direction only, and each arc has a given capacity. As shown in Fig. 2(a), the capacity of the arc connecting the source and A is 7 units, the capacity of the arc connecting the source and C is 3 units, and so on. With these given capacities, what is the maximum total flow that can be achieved from the source to the sink? It is quite apparent that we can regard this as a linear programming problem, with the variables being the flows in the arcs and the objective function being either the sum of the two flows going out of the source or the sum of the two flows going into the sink. Therefore, we can solve it by the regular techniques of linear programming. Unfortunately, the networks met in real life (e.g., railway networks, highway networks), are much more complex than the one in our example, with so many arcs (often thousands of them) that linear programming methods are not really practical.

There is a special approach, called the search for a flow-augmenting path, which makes a problem of this structure tractable. Consider the flow going through the network as shown in Fig. 2(b). The flow over the arc connecting source and node A is 4 items, the flow over the arc connecting



(a) Capacities of arcs



(b) Finding a flow-augmenting path

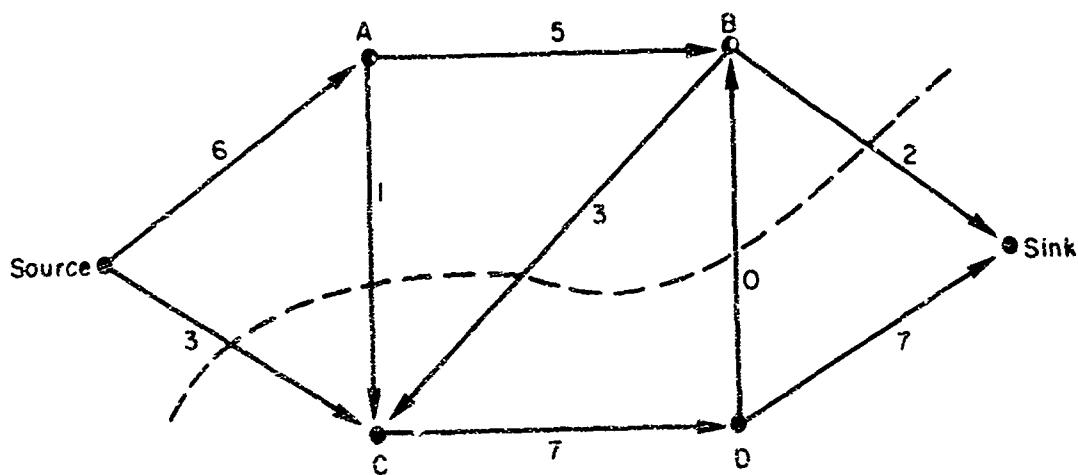


Fig. 2 - Flow in a simple network

nodes A and B is 3 items, and so on. This is a balanced flow, since at each node the sum of the income equals the outgo. Is this the maximum flow that can be achieved? Is this the most material that can be moved from the source to the sink? To answer that question, the straightforward and logical thing to do is look for a way to improve the flow, to look for what we call a flow-augmenting path. It is clear that we cannot ship anything more through the lower arc, from the source to node C—capacity is 3, flow is 3. However, through the upper arc, from the source to node A, we are shipping only four units, and the arc has a capacity of 7. We can tentatively add some quantity ϵ to this flow. At node A we now have some additional material ϵ which we want to ship out somehow. The arc to node B has flow 3 and capacity 6, so we can ship ϵ additional units to node B along it. We cannot ship it farther through either of the arcs directed outward (to node C and to the sink) from node B because they are both saturated. However, the other arc, coming into node B from node D, has a nonzero flow in it, and we can ship it from node B through this arc by reducing the flow through this arc by ϵ . We now have ϵ more material at node D which we could ship out by either of two ways—by reducing the flow from node C or increasing the flow to the sink. If we reduce the flow from node C, we will have an excess ϵ at C; and since we have already visited all C's neighbors, this is obviously not the way to go. Therefore, we increase the flow to the sink by ϵ ; and we see that we

have what is called a flow-augmenting path for the flow in this network. The path permits augmented flow, but by how much? How big can we make ϵ ? We can obviously make epsilon as big as 2. When we replace ϵ by 2 (Fig. 2(c)), the flow from the source to node A becomes 6; from A to B, 5; from B to D, 0; from D to the sink, 7. We can see that we now have a maximum flow through this network by considering the dashed line shown in Fig. 2(c) and examining the flows across this line in the direction from the source to the sink. The flow through each arc crossing this cut from source toward sink cannot be increased. Therefore, it is clear that there is no possible way of increasing the flow from the source to the sink beyond that shown in Fig. 2(c).

One of D. R. Fulkerson's great contributions to this subject was his discovery that a flow x in a network is maximal if, and only if, there is no flow-augmenting path with respect to x . That is to say, if the flow in a network is not maximal, a flow-augmenting path can always be found. Any nonmaximal flow whatsoever through a network can always be improved into an optimal flow by adding flow-augmenting paths. The search for a flow-augmenting path is a much quicker and simpler process than the solution of linear equations or finding feasible points with respect to linear inequalities.

The algorithms created and programmed for solving flows in networks problems are applicable to many areas

of operations research. (The February 1966 article by Fulkerson, which is cited in the bibliography, gives details and examples of such applications.)

GAME THEORY

As a simple example of game theory, consider a situation in which a boy wants to steal some fruit from a farmer who has an apple orchard and a watermelon patch. During any given night, the farmer can choose to guard either the watermelon patch or the apple orchard. The boy also has two choices: During any given night, he can try to steal either apples or watermelons. If the farmer is guarding the apples when the boy tries to steal apples, the boy will get nothing—he will get no payoff. But if the farmer is guarding the watermelons when the boy raids the apple orchard, the boy will get apples—he will get a payoff. If the boy likes watermelon three times as much as he likes apples, we can say that his payoffs from his raids will be 0 if he gets nothing, 1 if he gets apples, and 3 if he gets watermelons. Figure 3 shows the boy's payoffs under the four possible combinations of his choices of raids and the farmer's choices of guard posts.

What strategy should the boy use and what strategy should the farmer use? Let's draw a graph here (Fig. 4). Let's say that the farmer has a probability of guarding the apples which can range between 0 and 1. The boy can use one or the other of two strategies—Strategy I: raid the

Farmer's Strategies	
	I-Guard Apples II-Guard Watermelons
I-Steal Apples	0
Boy's Strategies	
II-Steal Watermelons	+1
	0

Fig. 3 - A zero-sum two-person game

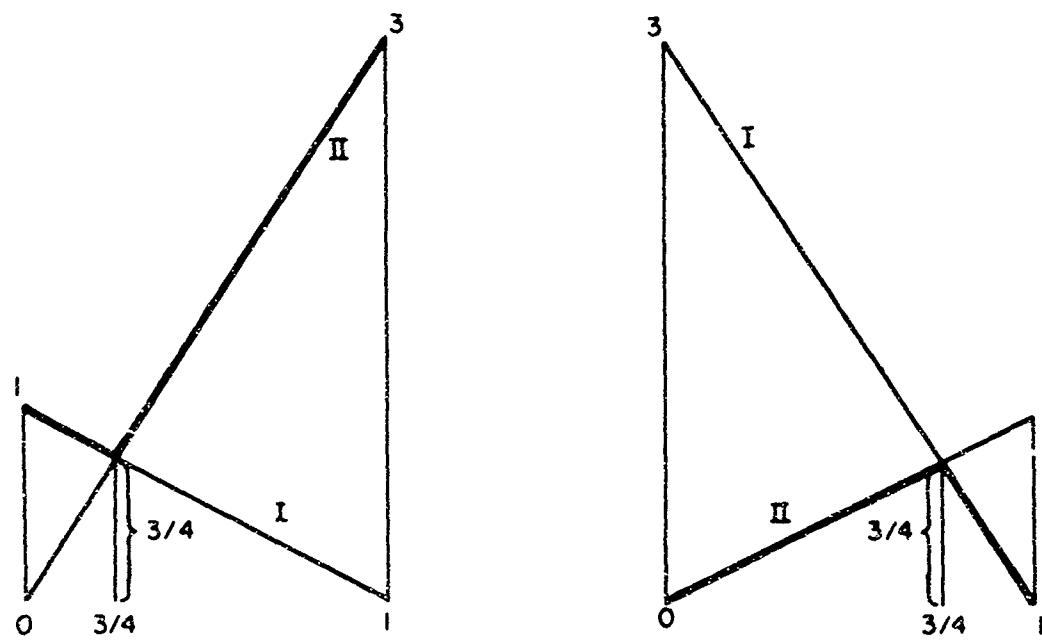


Fig. 4 - Choosing a strategy in a zero-sum two-person game

apple orchard; Strategy II: raid the watermelon patch. Following Strategy I, he gets 0 payoff if the farmer is guarding the apples and a payoff of 1 if the farmer is guarding the watermelons. If the farmer guards the apples with probability p and the watermelons with probability $(1 - p)$, the expected payoff to the boy (if he follows Strategy I) will be $0 \cdot p + 1 \cdot (1 - p)$. Following Strategy II, the boy will get a payoff of 3 if the farmer is guarding the apples and a 0 payoff if the farmer is guarding the watermelons. If the farmer guards the apples with probability p and the watermelons with probability $(1 - p)$, the expected payoff to the boy (if he follows Strategy II) will be $3 \cdot p + 0 \cdot (1 - p)$.

The boy wants to make his payoff as large as possible, of course, so he will make a guess as to the probability that the farmer will be guarding the apples, then follow the strategy that has the greater expected payoff. The crossover point between the two strategies occurs at $p = 1/4$ and the value that the boy achieves here is $3/4$. If the probability that the farmer will be guarding the apples is less than $1/4$, the boy should go for the apples. If that probability is greater than $1/4$, the boy should go for the watermelons. If the probability is exactly $1/4$ for a guard on the apples and $3/4$ for a guard on the watermelons, it makes no difference which strategy the boy follows. If the farmer mixes his strategies in this way, he can make

sure that the boy will not do any better than 3/4 expected payoff.

Let's consider this game from the farmer's standpoint. He also has two strategies—Strategy I: guard the apples and Strategy II: guard the watermelons. He has to make a judgment about where the boy is probably going to strike. If the farmer feels that there is a 0 probability of an apple raid, he will be permitting a payoff of 3 if he follows his Strategy I and a payoff of 0 if he follows his Strategy II. Similarly, if he feels that the boy is certainly going to hit the apples, a 0 payoff will be permitted if the farmer guards the apples and a payoff of 1 if the farmer guards the watermelons. The expected payoff at intermediate probabilities is found, as before, by linear interpolation. The crossover point here occurs at 3/4, and the payoff to the boy is 3/4. The farmer is trying to minimize the boy's payoff. Therefore, if he feels that there is less than 3/4 probability that the boy will hit the apples, he should guard the watermelons; but if that probability is greater than 3/4, he should guard the apples. Note that if the boy follows the strategy of hitting the apples 3/4 of the time and the watermelons 1/4 of the time, he will always have expected payoff of 3/4.

It is very interesting that the boy can guarantee himself a payoff of 3/4 by following this division of strategy and that the farmer can be sure that "the boy does not make anything better than a payoff of 3/4 by following

his strategy of $1/4$ guarding apples, $3/4$ guarding watermelons. The maximum expected payoff the boy can guarantee to himself is exactly the same as the smallest expected payoff the farmer can limit him to. This result is valid beyond these simple 2 by 2 matrix games; it holds true for an n by m matrix, for example, one with a hundred different strategies for the farmer and fifty different strategies for the boy. However many strategies were available to each, there would be some probability distribution for the farmer that would guarantee limiting the boy to a payoff of a certain expected value and some probability distribution for the boy that would insure him an expected payoff of no less than exactly the same value. This is the famous minimax theorem of Von Neumann, the fundamental theorem of zero-sum two-person games. This very striking result, which puts a value for each player on playing a game of this sort, makes this kind of game theory applicable to a great many military problems (e.g., ballistic missile defense, balancing the hardness of missile silos between the silos themselves and the control centers, balancing defense against missiles and bombers).

However, a great many of the conflict situations that arise in real life are not really zero-sum games. Let's suppose that our boy, encouraged by stealing apples, grows up and becomes an adult delinquent. He and his buddy are arrested and the district attorney takes our boy off to

interrogate him separately from his buddy. "Now look," he says, "if you don't talk, I'm going to hold you for a week on vagrancy. But if you do talk and implicate your friend [in other words, if you cop out and your buddy doesn't, I'll let you go right away. But, of course, your friend will be in jail for ten weeks." Simultaneously, the assistant district attorney is saying exactly the same thing to Boy B. The payoffs in this game are shown in Fig. 5. If both of them refuse to talk, they will each spend a week in jail. If B cops out and A does not cop out, A goes to jail for ten weeks and B goes free at once, with a symmetric pattern for the inverse. If they both cop out, they both go to jail for, say, nine weeks.

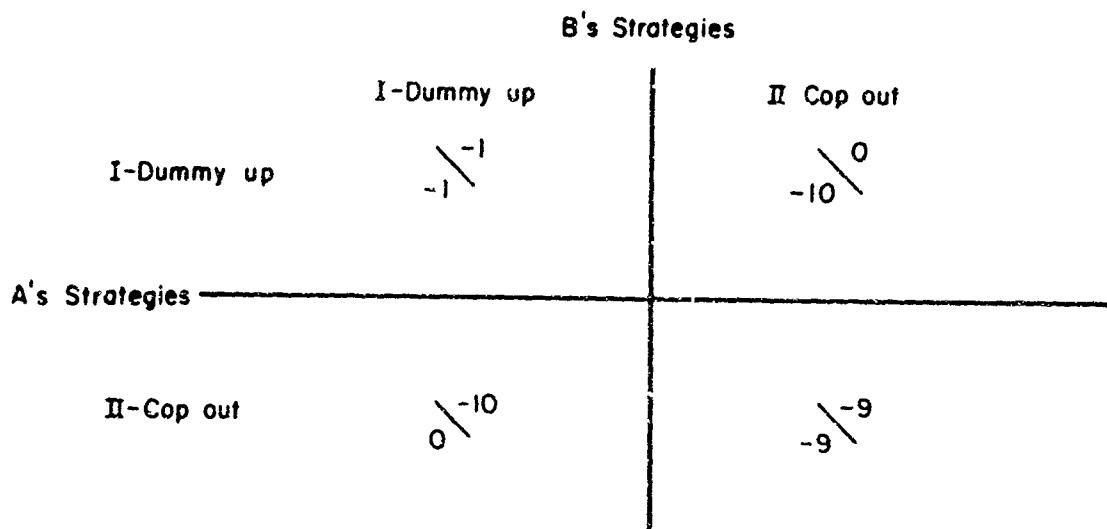


Fig. 5 - A nonzero-sum two-person game

In game theory, this classic problem is known as "The Prisoner's Dilemma," and it really defies simple analysis, since so much depends on the conditions in the outside world. If one of the outside conditions is a syndicate that kills a squealer very dead within 24 hours after he leaves the police station, certainly neither of the boys will cop out. They will clam up as a life or death matter, spend a week in jail, and then go free. Since either or both of them could get a nine- or ten-week sentence, probably they are both better off if there is a syndicate than if there is not a syndicate—perhaps I should say (I don't want to get into moral overtones) that they are both better off if they believe there is a syndicate.

This is an illustration of one of the difficulties which game theory gets into beyond the zero-sum two-person area. The prisoner's dilemma illustrates the difficulties that arise when the game is no longer zero-sum, when one player's gain does not always mean a symmetric loss to the other player. In other words, when players have common interests or common fears in a situation, the minimax theorem no longer applies and the simple solution concept associated with it no longer applies.

Another kind of difficulty arises with a game that is zero-sum but involves more than two players. For an example, consider three players—A, B, and C—who have agreed to participate in a shooting match. Player A has

a balloon, Player B has a balloon, and Player C has a balloon; and these balloons are equidistant from one another. The players decide by lot who will have the first shot. The player who has the first shot takes his pistol and fires at the balloon of one of the other players. If he misses the balloon, they draw lots again. If he hits the balloon, the balloon pops, and the player who is holding it is out of the game.

The probability that A will hit a balloon he shoots at is a , the probability for B is b , and the probability for C is c . If C's balloon has been popped and C is out of the game, so that A is playing against B, it is an easy matter to compute P_{AB} , the probability that A will win against B. There is probability $1/2$ that A will get the shot on the drawing and there is probability a that he will win with that shot, so there is probability $a/2$ that A will simply win at once. However, there is a probability $(1 - a)/2$ that he will get the first shot, but will miss. Similarly, there is a probability $(1 - b)/2$ that B will get the first shot, but will miss. If either of these two events happen, we are back where we started, and A has probability P_{AB} of winning against B. So we have a simple linear equation, and there is no need to go through the mechanics of solving it here:

$$P_{AB} = \frac{a}{2} + \left(\frac{1-a}{2} + \frac{1-b}{2} \right) P_{AB}.$$

The solution is simply that the probability of A winning against B is $a/(a + b)$, and similarly for any other pair of players.

When only two players are playing, there can be no real strategy. If you win the toss, you shoot your gun and hope you hit the balloon. You really have no choice—you just try to do the best you can on each shot and you make no strategic decisions in the matter. When there are three players, you do make a strategic decision: You decide which player's balloon you will shoot at.

Assume that all players' probabilities of hitting are known to all the players, and that A's probability of hitting is greater than B's, which in turn is greater than C's. The problem of strategy is thus apparently simplified because obviously the balloon of the more dangerous opponent is the preferred target, and a player's probability of hitting is the same for it as for the other. It makes sense—it is good strategy to eliminate the more dangerous opponent first. We can easily compute P_A , the probability that A will win when all players follow this sensible strategy. There's $1/3$ chance that A will get the first shot. He has probability a of hitting when he shoots. He will obviously shoot at the balloon of B, the more dangerous opponent. If he hits B's balloon, he will consider his chances of winning against C. There are alternatives. Player A may get the first shot and miss. If B gets the first shot, he will of course shoot at A's balloon; and if he hits it, A will be out of

the game—so that contingency can be ignored because it contributes nothing to A's chances of winning. If C gets the first shot, C will shoot at A's balloon and if he hits, that contingency can also be ignored. But if A gets the first shot and misses, or B gets the first shot and misses, or C gets the first shot and misses, the situation is the same as it was originally. This gives the following equation:

$$P_A = \frac{a}{3}P_{AC} + \left(\frac{1-a}{3} + \frac{1-b}{3} + \frac{1-c}{3}\right)P_A$$

A similar line of reasoning gives the following equations for P_B and P_C , the probabilities that B and C win:

$$P_B = \left(\frac{b}{3} + \frac{c}{3}\right)P_{BC} + \left(\frac{1-a}{3} + \frac{1-b}{3} + \frac{1-c}{3}\right)P_B$$

$$P_C = \left(\frac{b}{3} + \frac{c}{3}\right)P_{CA} + \left(\frac{1-a}{3} + \frac{1-b}{3} + \frac{1-c}{3}\right)P_C$$

This system of three linear equations in three unknowns (the quantities P_{AC} , P_{BC} , and P_{CA} are known) can be simply solved to yield the following formula: A's probability of winning will be

$$P_A = \frac{a^2}{(a+b+c)(a+c)}$$

B's probability of winning will be simply

$$P_B = \frac{b}{(a+b+c)}$$

and C's probability of winning will be

$$P_C = \frac{c(2a + c)}{(a + b + c)(a + c)}$$

Everything seems highly logical. All is right with the world. When all players follow a reasonable strategy, this is the outcome. But let's try plugging in some specific numbers and see what answers we get. When a has the value .8, and b has the value .6, and c has the value .4, A 's probability of winning is .296, B 's probability of winning is .333, and C 's probability of winning is .370. In other words, the worst shot has the best chance of winning because nobody is afraid of him. Nobody is shooting at his balloon. A and B are gunning for each other, and C is just sitting back there taking a pop at the big kids when he gets a chance.

If A and B realize this, they can get together and say, "Well, look, it's foolish for us to strike against each other when C here will just clean up after us. So why don't we make a deal that first we wipe out C and then we fight each other?" If A and B go after C first, of course there will be a different system of linear equations to solve. Solving these new equations and again substituting .8, .6, and .4 for a , b , and c , yield the following results: A now has a good chance of winning, .444; B has a slightly better chance, .465, because C will still shoot at A 's balloon first before shooting it out with B ; and C has only 9 chances out of 100 of emerging victorious from this gang-up situation.

Of course, it is possible that the game is structured in such a way that A and B cannot reach this kind of agreement, although it is hard to see how they could be prevented from reaching it if they really understood the game in advance. If collusion with B is not possible, A could use a strategy that would give him at least as good a win probability as that resulting from collusion, provided he could convince C that he really is a very hard-nosed, "irrational" player. A at the beginning of the game should say that generally he will follow sensible strategy and go after B when he gets a chance, but if C shoots at A's balloon, A will retaliate and shoot at C's balloon. (This might be called a deterrent strategy.) It might seem that A would lessen his chances of winning because he would be attacking his less dangerous opponent. But if his threat is believed and he thus succeeds in deterring C from attacking him, so that if C drew first shot he would aim at B's balloon rather than A's, there will be a different system of equations. Solving them and making the substitutions, we find winning probability .444 for A, probability .200 of a win for B, and probability .356 of a win for C.

This example illustrates the fact that multi-person games cannot be properly discussed until adequate information is provided about the social climate, about the possibilities for communication, for compensation, and for commitment and trust. Given this information, one can proceed to the formulation of a suitable solution concept; and different

solution concepts from multi-person games are appropriate for different situations. It has been found that many of the large games, such as a large market, are actually easier to solve than the smaller games. In other words, many of the problems based on real-world situations are easier to deal with than these created examples.

The theory of multi-person games gives new justification for many classical economical theories and provides new insights into many economic situations. Lloyd Shapley of the RAND staff has collaborated with Martin Shubik, an economist from Yale University, in writing Competition, Welfare, and the Theory, of Games, in which multi-person game concepts are used to give some insight about present-day economic problems.

Recently the Committee on the Undergraduate Program in Mathematics of the Mathematical Association of America published its model curriculum in applied mathematics. This curriculum comprises two options: the physical sciences option and the optimization option. The latter option comprises twelve named topics. Three of the twelve are topics I have discussed today: dynamic programming, linear programming, and the theory of games. Two others are topics I have not discussed, but which are areas in which RAND work is cited by the committee: scheduling problems and nonlinear programming. The fact that so much of the undergraduate program consists of mathematics which was essentially created by the small band of mathematicians

at RAND is a fact which should give us great satisfaction.
It shows that our eye has been on the ball in the past.
Let's hope it stays on the ball in the future!

BIBLIOGRAPHY

Linear Programming

Dantzig, G. B., Linear Programming and Extensions, Princeton University Press, Princeton, New Jersey, 1962 (\$11.50).

Dynamic Programming

Bellman, Richard, Dynamic Programming, Princeton University Press, Princeton, New Jersey, 1957 (\$8.50).

Bellman, Richard, and S. Dreyfus, Applied Dynamic Programming, Princeton University Press, Princeton New Jersey, 1962 (\$8.50).

Flows in Networks

Ford, L. R., Jr., and D. R. Fulkerson, Flows in Networks, Princeton University Press, Princeton, New Jersey, 1962 (\$6.00).

Fulkerson, D. R., "Flow Networks and Combinatorial Operations Research," American Mathematical Monthly, February 1966.

Game Theory

Dresher, Melvin, Games of Strategy: Theory and Applications, Prentice-Hall, 1961 (\$9.50).

Shapley, L. S., and M. Shubik, Competition, Welfare and the Theory of Games, to be published.